

CHAPTER 8 IP SECURITY

After reading this chapter and completing the exercises, you will be able to:

- ◆ Describe the features and benefits of the IP Security protocol
- ◆ Describe the two modes of operation for IP Security: transport and tunnel
- ◆ Describe the IP Security authentication and architecture
- ◆ Configure IP Security for transport mode on a Windows 2000 server
- ◆ Configure IP Security for tunnel mode on a Windows 2000 server
- ◆ Customize IP Security policies and rules
- ◆ Manage and monitor IP Security

Security has always been an important issue in computer networking, and the wealth of security options the modern administrator can choose from proves this. The many different forms of networking security include the authentication of users when they log on to the network; the restriction of physical access to computers and networking equipment; the authentication of computers as they pass data to other computers; and the protection of the actual data being sent across the networks. IP Security (IPSec) falls into this last category. IPSec is an extension of the IP protocol that provides point-to-point encryption of data being sent between two computers on an IP-based network.

This chapter begins with an overview of the benefits, features, and operations of IPSec. From there, it moves on to cover the actual implementation, configuration, and management of this promising security protocol.

IPSEC OVERVIEW

IPSec is an extension to the familiar **Internet Protocol (IP)** that is responsible for routing data packets on a TCP/IP-based network. Actually, IPSec is not a single protocol, but a suite of protocols designed to work together to secure data being sent between two computers on a network. Like IP, IPSec works at the Network layer. This means that higher-level protocols and applications in the TCP/IP protocol suite, like FTP, can ignore the encryption process. They carry out their functions normally, passing data down the protocol layers, unaware of whether that data is eventually encrypted or not.

Any Windows 2000 computer may act as an **IPSec client** or an **IPSec server**. The IPSec client is the computer that initiates the IPSec connection, and the IPSec server is the one that receives it. Nothing special about IPSec configuration makes a computer an IPSec client or server; it just depends on which computer makes the initial connection attempt.

This overview examines what functions the IPSec protocol actually performs and the benefits that it provides. It also looks at the individual protocol components that make up IPSec and how they all work together to send encrypted data from one computer to another.

What IPSec Does

IPSec is a framework of open standards developed by the IPSec working group of the Internet Engineering Task Force (IETF). The framework provides a way to ensure transfer of encrypted data over IP-based networks. The Microsoft Windows 2000 implementation of IPSec is based on those standards.

IPSec provides two basic services. The first is a way for two computers to determine whether they trust one another. This is referred to as **authentication**. The second service is to provide a reasonable way to **encrypt** data on one end of the connection and **decrypt** it on the other end. IPSec is often called an end-to-end security measure, meaning that only the sending and receiving computers must know about IPSec. The medium over which the data travels is assumed to be insecure, and the IPSec process requires no other computers on the network to be involved. Also, routers that forward packets of data between networks also do not need to know anything about IPSec.

Features of IPSec

Since it is enabled at the Networking level, the greatest benefit of IPSec is that it is completely transparent to users, applications, and protocols above and below the Networking layer. The following list describes some of the additional features offered by the Windows 2000 implementation of IPSec:

- IPSec uses the Windows 2000 domain as a trust model. By default, **IPSec policies** use the default Windows authentication method (**Kerberos V5**) to validate communicating computers. IPSec policies can also be configured to use **public key certificates** or **pre-shared keys** for authentication.

- **IPSec policies**, sets of rules assigned to clients that define how those clients use IPSec, are assigned centrally through the **Active Directory Group Policy** feature. This chapter covers IPSec policies in detail later.
- All packets are encrypted using time-specific information so that they cannot be captured and played back later in an attempt to crack the encryption code.
- Long key lengths and dynamic changes of keying are used during ongoing communications for added security. (You can learn more about the basics of encryption in Chapter 10.)
- Private network users can connect using secure end-to-end links with any trusted domain in the enterprise.
- Remote users and private network users can connect using secure end-to-end links based on IP addresses.

Modes of Operation

IPSec can operate in two different modes, depending on the scope of the communication. These two modes of operation are transport mode and tunnel mode.

Transport Mode

When used to secure communication between two specific clients, such as two computers on the same LAN, IPSec operates in **transport mode**. The two endpoints of communication are the two computers, and both must have IPSec configured. For this mode to work, both computers must use the TCP/IP protocol.

Tunnel Mode

IPSec can also be used to secure communication that passes through a transit network such as the Internet. In this case, called **tunnel mode**, the two communicating computers do not use IPSec themselves. Instead, the gateways connecting each client's LAN to the transit network create a virtual tunnel that uses the IPSec protocol to secure all communication that passes through it. Communication from the clients themselves is encapsulated in the tunnel protocol headers, encrypted, and passed through the tunnel. The gateway at the other end decrypts the packet, removes the tunnel protocol header, and sends the packet to the destination computer.

Tunnels can be created using only the IPSec protocol or by combining IPSec with the Layer 2 Tunneling Protocol (L2TP) to establish a Virtual Private Network (VPN) connection. In this case, L2TP, rather than IPSec, actually creates the tunnel. For more information on L2TP and VPNs, see Chapter 6.

In tunnel mode the actual clients are not involved in IPSec communication. This frees the clients to use other networking protocols such as IPX/SPX and AppleTalk; they are not restricted to TCP/IP.

IPSec Authentication

As mentioned previously, IPSec supports three different types of authentication:

- **Kerberos** is the default authentication system used by Windows 2000. An open standard, it is thus widely supported by other operating systems, as well.
- **Certificates** are provided by a certificate authority. Each end of the IPSec connection uses the other end's public certificate for authentication. This model provides good security but also requires that a certificate server be accessible for the distribution of certificates. Chapter 10 covers this topic in detail.
- **Pre-shared keys** are simply passwords entered into each computer. As long as both computers are configured with the same pre-shared key, they trust one another. While the pre-shared key itself is never transmitted between the clients, it is stored in the Active Directory in an unencrypted format. For this reason, pre-shared keys are considered less secure than the other available forms of authentication.

IPSec Architecture

Once the two communicating clients authenticate one another, they are ready to begin encrypting and sharing data. This section first introduces the various components that are a part of this process and then describes the process itself.

IPSec Components

IPSec is implemented using a number of different components. The following sections introduce these.

IPSec Policy Agent Service The **IPSec policy agent service** resides on each Windows 2000 computer that is configured with IPSec. It starts automatically when the computer starts and performs several tasks at specified intervals. The policy agent is responsible for retrieving the computer's assigned IPSec policy from the Active Directory. If it cannot connect to the Active Directory, or if it finds no policy there, it tries to retrieve the policy from the computer's registry. If it finds no policy, IPSec cannot continue. If it does find a policy, the policy agent sends the information in the policy to the ISAKMP/Oakley Service.

ISAKMP/Oakley Service A key management service, the **ISAKMP/Oakley Service** also resides on each Windows 2000 computer involved in IPSec communication. Before two computers attempting a connection can send any data, they must first establish a security association. This **security association** defines the common security mechanisms, such as keys, that the two computers use to create the IPSec connection. The ISAKMP/Oakley Service is also responsible for generating the keys used to encrypt and decrypt the data sent over the IPSec connection.

IPSec Driver The **IPSec driver** also resides on each Windows 2000 computer involved in IPSec communication. The IPSec driver starts when the policy agent starts. It watches all IP

datagrams for a match with a filter list configured in the computer's security policy. Filters are used to define what computers can and cannot establish connections with other computers. If it finds a filter match, the IPSec driver uses the keys created by the ISAKMP/Oakley Service to encrypt the data and send it over the IPSec connection. The IPSec driver on the receiving computer decrypts the data.



To force the IPSec driver to restart, you can restart the IPSec Policy Agent using the Services console in the Administrative Tools program group.

The IPSec Process

Now that you've been introduced to the major components involved in IPSec communication, here's how they all work together:

1. An application on one computer (call it Host 1) sends data to another computer (call it Host 2).
2. The data passes down through the networking layers of Host 1, where it is fragmented and shaped into packets to be sent over the network.
3. When the data reaches the networking level and is ready for routing by the Internet Protocol, the IPSec driver for Host 1 notifies the ISAKMP/Oakley Service that an IPSec connection is needed.
4. The ISAKMP/Oakley Services on both computers establish a security association and generate a shared key.
5. The ISAKMP/Oakley Services on both computers transfer the shared key to the IPSec drivers on those hosts. Now, the IPSec drivers on both computers have the same shared key.
6. The IPSec driver on Host 1 uses the key to encrypt the data and then sends the data to Host 2.
7. The IPSec driver on Host 2 receives the data and uses the shared key to decrypt it.
8. The IPSec driver passes the data up to the next networking layer.
9. When the data works its way up to the top layer, the application on Host 2 receives the data and never knows it was encrypted.

INSTALLING IPSEC

All IPSec components are installed by default when you install Windows 2000. All you really need to do to enable IPSec is to create a custom console using the Microsoft Management Console that includes the IP Security Policy Management snap-in (called the IPSec snap-in from now on) and then use the snap-in to assign policies and filters. Hands-on Project 8-1 at the end of the chapter outlines the actual steps for creating a console and assigning a policy.

During the creation of the snap-in console, there is really only one option to which you have to give some thought. You must choose whether you want the snap-in to manage local IPsec policy, the default policy for your computer's domain, the default policy for another domain, or the local policy on another computer. See Figure 8-1.

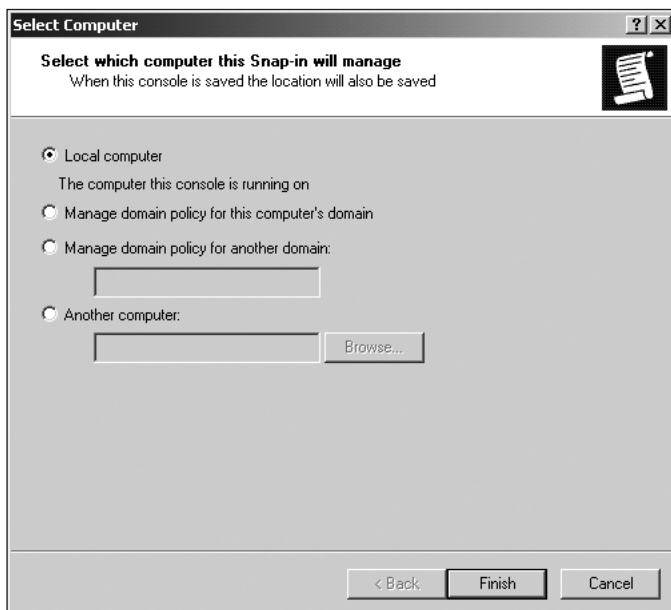


Figure 8-1 Enabling IPsec



While the real power of IPsec lies in configuring group policies in the Active Directory, a discussion of group policy management is really outside the scope of this book and, for the most part, the scope of the certification exam, as well. Instead, this chapter focuses on customizing and controlling the IPsec settings themselves. While you can manage policies at a variety of levels, you always use the IPsec snap-in to manage them. Furthermore, the management of policies at the local level and at the Active Directory level uses essentially the same techniques.

CONFIGURING IPSEC

You can modify the existing IPsec policies to suit your needs, create policies of your own, or both. You create policies using a wizard that steps you through the process of configuring the policy. You manage the policies you create using property pages, just like you manage objects in other snap-ins you work with. For each policy, property pages allow you to configure general settings and a set of rules under which the policy operates. The Rules page lets you tie general filters you create to filter actions taken when that filter is in effect. This

section covers all of this creation and configuration. First, though, it may be helpful to get our bearings straight with regard to the IPsec snap-in, shown in Figure 8-2.

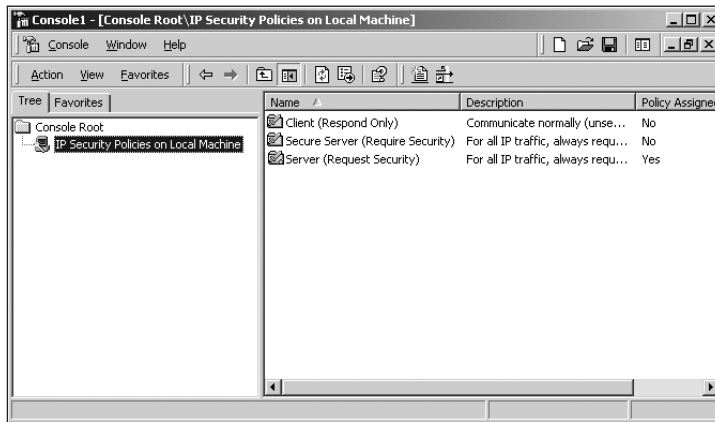


Figure 8-2 IPsec snap-in

Selecting the IP Security Policies on Local Machine object displays a list of available policies in the right pane. The list describes each policy and also tells whether the policy is assigned, or functional. You can right-click any existing policy and choose the Properties command from the shortcut menu to configure the policy. The same shortcut menu also contains commands for assigning or unassigning the policy, depending on its current state.

As shown in Figure 8-3, right-clicking the IP Security Policies on Local Machine object reveals that there are no properties to configure for that object, but that you can perform a number of tasks at this level.

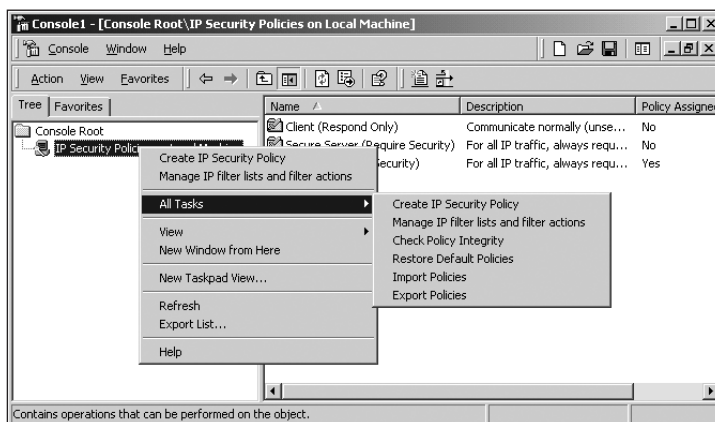


Figure 8-3 IPsec tasks for a local machine

Your main tasks are creating a new policy and managing the list of filters and filter actions available for use in the rules you create for policies. Other tasks include checking policy integrity, restoring default policies, and importing/exporting policies for use in other IPSec snap-ins. The following sections cover all these tasks.

Creating a New Policy

Creating a new policy is a wizard-based process that you start by right-clicking the IP Security Policies on Local Machine object and choosing New IP Security Policy from the shortcut menu. The first couple of the wizard's pages simply inform you of the actions the wizard performs and ask you to name the new policy—simple enough. The next page, shown in Figure 8-4, asks whether you want to enable the default response rule for the policy.

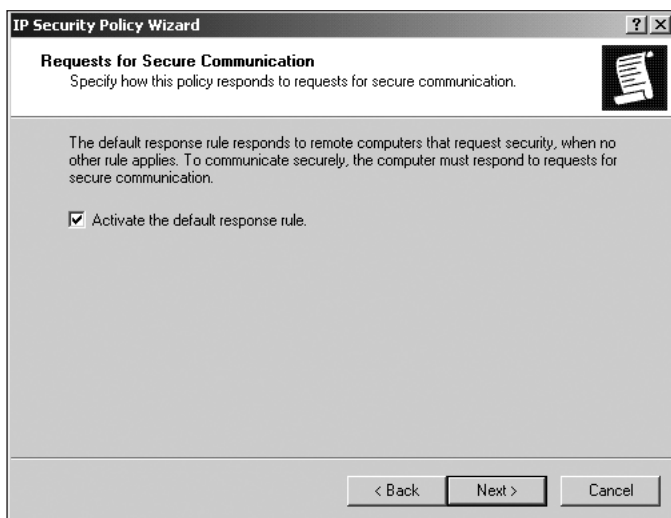


Figure 8-4 Default response rule for a policy gives permission to accept a connection

This default rule basically permits the local computer to accept an IPSec connection from anyone requesting one. Unless you customize the default rule (something discussed a bit later in the chapter), it's probably a good idea not to enable it. It's better to set policies that allow only known hosts to connect. If you choose not to use the default rule, the wizard finishes right then and creates the new policy.

If you do enable the default rule, the wizard next asks you to configure an authentication method for the rule, as shown in Figure 8-5. You can choose any of the three authentication methods (Kerberos, certificates, and pre-shared keys) discussed earlier in the chapter.

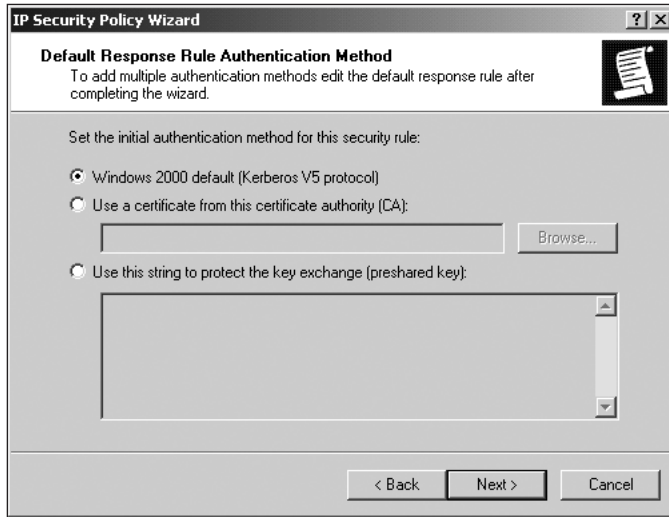


Figure 8-5 Choosing an authentication method for a new rule

Configuring a Policy

Once you create a policy, you can configure it by right-clicking the policy and choosing Properties from the shortcut menu. A policy holds only two property pages, General and Rules, which the following sections discuss. Hands-on Project 8-2 at the end of the chapter gives you a chance to practice configuring a policy.

General Properties

The General page for a policy, shown in Figure 8-6, lets you change the name of the policy and the description. Even though these items only appear in the IPSec snap-in, it's a good idea to create descriptive names that help you identify the policies. The Check for policy changes every *x* minutes field lets you change the interval at which clients that use this policy check to see if the policy has been updated.

The Advanced button opens the Key Exchange Settings dialog box shown in Figure 8-7. This dialog box lets you control how often the policy requires the communicating computers to regenerate new keys. The default is after about 480 minutes (eight hours), but you can change this to any value you like or configure it to require new keys after a set number of sessions. While the default setting usually works fine, regenerating new keys more often is more secure. You must strike a balance between your need for security and the time consumed generating new keys.

The Methods button displays a list of security methods used to exchange the keys. The Master key Perfect Forward Secrecy option lets you prohibit the reuse of keying material or keys. If you select this option, you may only set the regeneration of keys to occur at timed intervals. The sessions field becomes unavailable.



Figure 8-6 General page for a policy

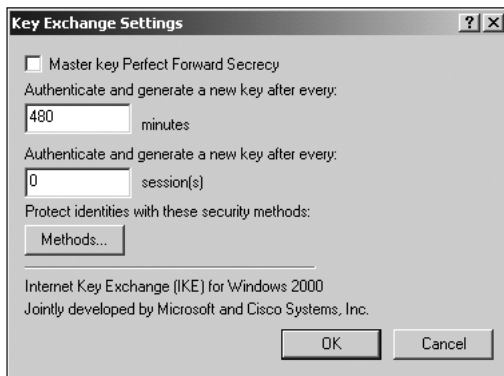


Figure 8-7 Key Exchange Setting dialog box

Rules Properties

You use the Rules page of a policy, shown in Figure 8-8, to define the rules included with that policy. Each rule listed includes the following entries:

- check box to the left of the rule, which specifies whether the rule is actually turned on or off
- Filter List, a list of filters that defines the connections to which a particular rule applies

- Filter Action, determines what happens when a connection meets the criteria set by a filter list
- Authentication Method that the rule uses
- Tunnel Setting, covered later in this chapter
- Connection Type, defines the type of connection to which the rule applies

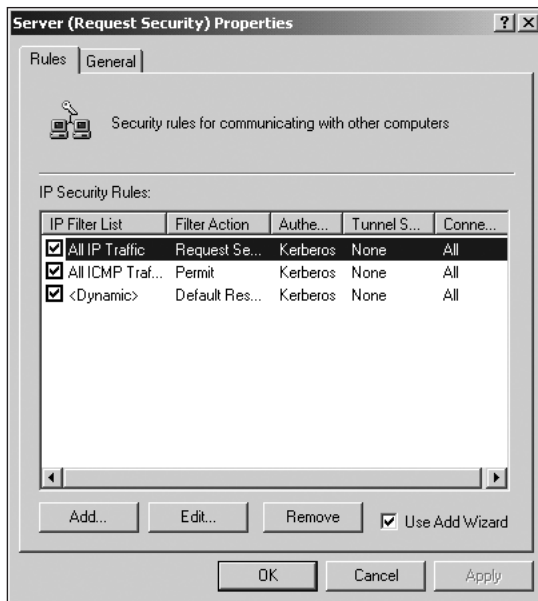


Figure 8-8 Rules page of a policy

You can use the Add button to create a new rule, the Edit button to open the property pages for a rule so that you can configure it, and the Remove button to delete a rule. When you add a new rule, one of two things happens:

- If the Use Add Wizard option is not enabled, the property pages for the new rule open and you can configure it directly.
- If the Use Add Wizard option is enabled, a wizard steps you through the configuration of the rule. Since the wizard really just asks you questions and then fills in the parameters on the property pages for you, we're just going to cover the property pages themselves rather than the wizard. Once you understand the property pages, the wizard will be easy to use. The following sections describe each of the property pages available for a rule.

IP Filter List Properties The IP Filter page, shown in Figure 8-9, shows all of the filter lists associated with the policy. The list includes all filter lists available on the server, and you simply select one to associate with the rule. You can create new filter lists here using the Add

button, this is a topic discussed later in the section, “Managing Filter Lists and Actions.” The reason we do not recommend constructing filter lists at the individual policy level is that filter lists apply to all policies and are better managed at that level.

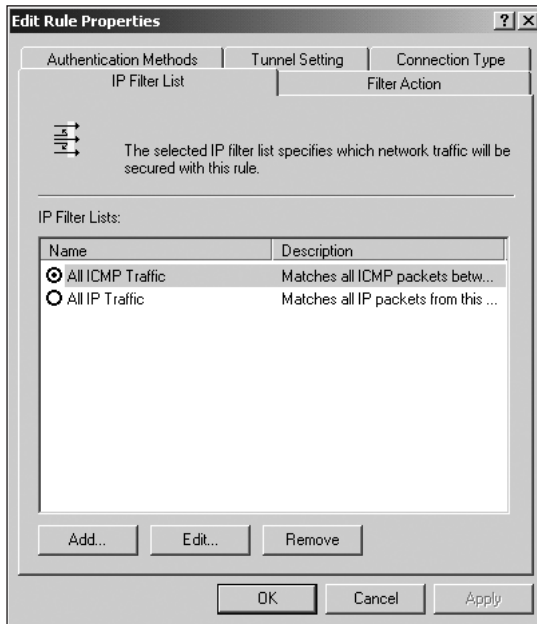


Figure 8-9 IP Filter page of a rule

Filter Action Properties The Filter Action page, shown in Figure 8-10, shows all of the filter actions associated with the policy. This list shows all filter actions available on the server, and you can select any one of them to associate with the rule. As with the Filter Lists page, you can create new filter actions from this page, but we find it more appropriate to create them at the level of the policy instead.

Authentication Methods Properties The Authentication Methods page, shown in Figure 8-11, lets you define one or more authentication methods to use with the rule. If you select more than one, IPSec attempts to use them in the order that they appear on the list. Earlier sections of this chapter discuss the three authentication methods, Kerberos, certificates, and pre-shared keys.

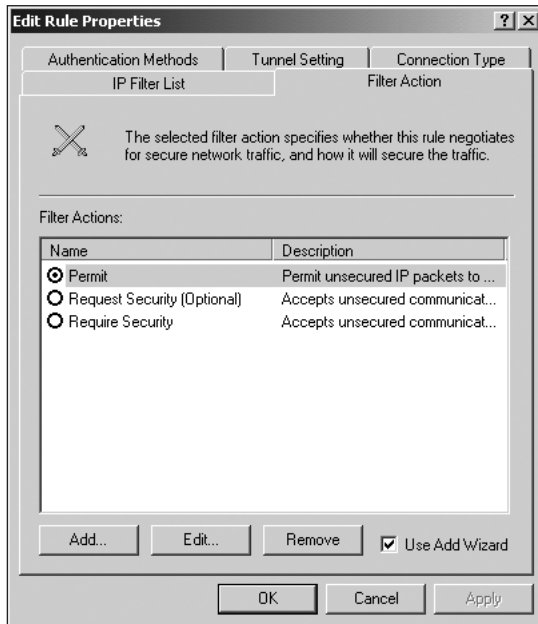


Figure 8-10 Filter Action page of a rule

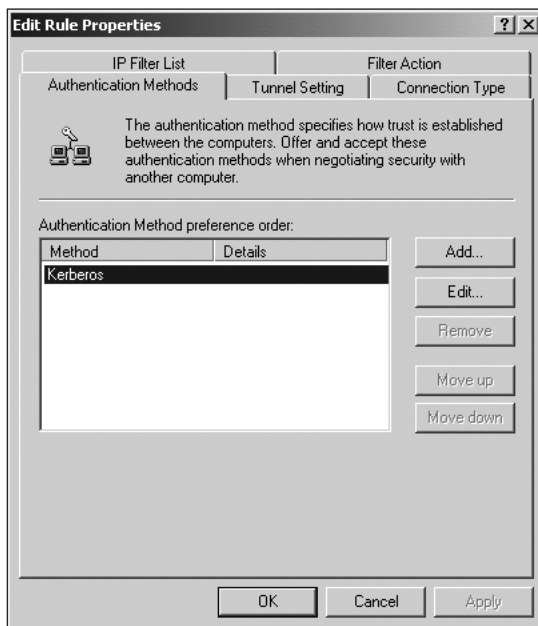


Figure 8-11 Authentication Methods page of a rule

Tunnel Setting Properties You use the Tunnel Setting page, shown in Figure 8-12, to specify whether or not a connection is tunneled. This means that you specify whether connections are tunneled on a per-rule basis. To enable tunneling, select the tunnel option and enter an IP address to serve as the tunnel endpoint.

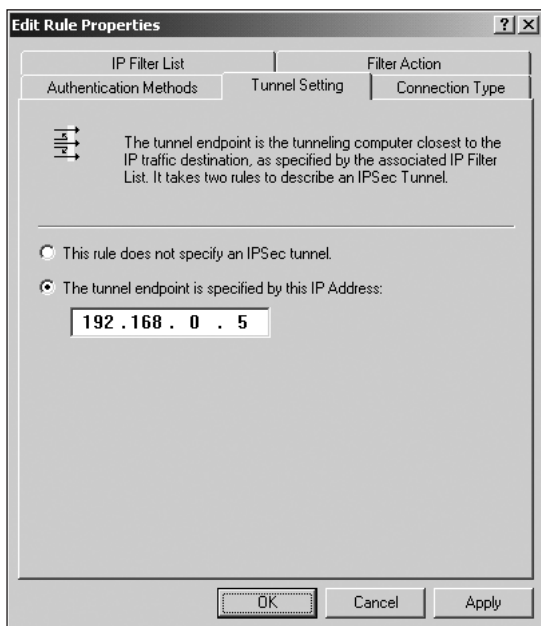


Figure 8-12 Tunnel Setting page of a rule

Of course, setting up a tunnel requires some additional steps. To properly construct a tunnel, you need two tunnel rules on both ends of the tunnel (one for inbound traffic and one for outbound traffic) with the appropriate filter lists and filter actions in place. You need to configure each end as follows:

- Configure an outgoing rule with a filter list that specifies the other end of the tunnel as the tunnel endpoint.
- Configure an incoming rule with a filter for incoming traffic from any subnet from the remote end of the tunnel.

Hands-on Project 8-4 outlines the steps involved in setting up one end of a tunnel.

Connection Type Properties The Connection Type page, shown in Figure 8-13, lets you specify the kind of connections to which the rule applies. Your choices are to have the connection apply to LANs only, remote access connections only, or all network connections (both LAN and remote access).

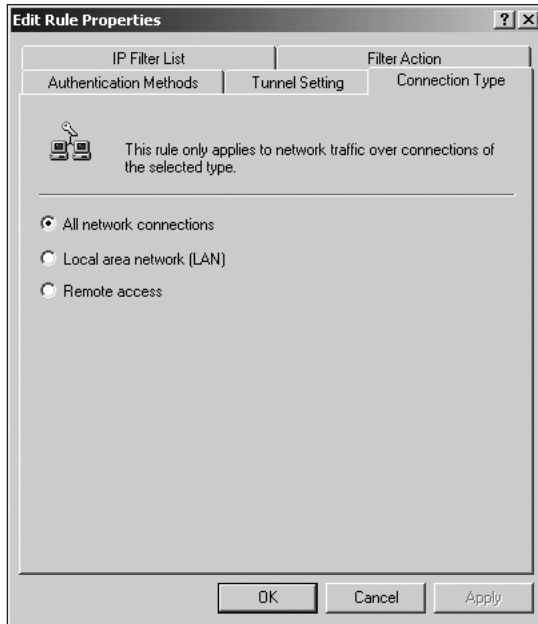


Figure 8-13 Connection Type page of a rule

Managing Filter Lists and Actions

The previous section explained how you can configure filter lists and actions from the Edit Rule Properties dialog box. However, since filter lists and actions are available to all policies creating and managing them at a higher level makes more sense. To do this, right-click the IP Security Policies on Local Machine object and choose the Manage IP Filter Lists and Filter Actions command from the shortcut menu. This opens a dialog box with two pages, Manage IP Filter Lists and Manage Filter Actions, which show you exactly what this dialog box is used for. The following sections discuss each of these pages.

Managing IP Filter Lists

You use the Managing IP Filter Lists page, shown in Figure 8-14, to manage filter lists available to all policies. This list simply shows the filters available for your policies and has nothing to do with the application of those filter lists.

To edit an existing filter list, select the list and click the Edit button. This opens a dialog box similar to the one shown in Figure 8-15. You can also click the Add button to open a blank version of the same dialog box and add a new filter. This dialog box shows a name and description for the filter list, as well as the actual filters in that list.

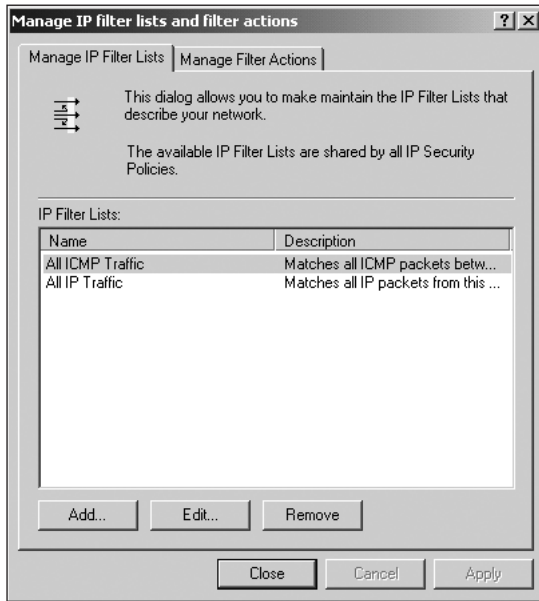


Figure 8-14 Managing IP Filter Lists

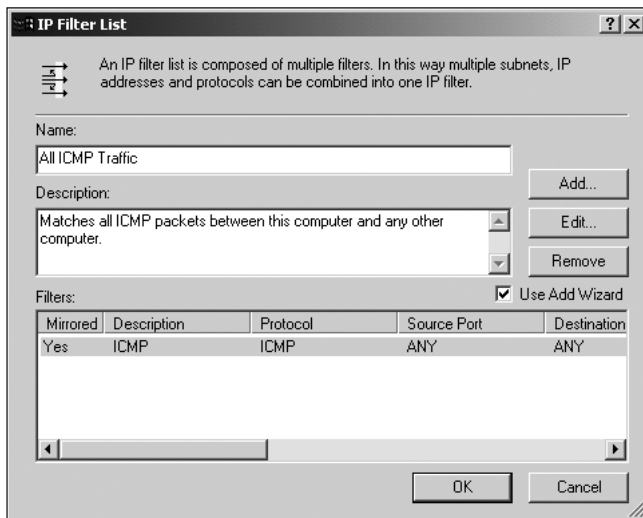


Figure 8-15 Editing filters on a filter list

The same kind of logic applies to creating and configuring individual filters that applies to creating rules. As you learned earlier, selecting a filter and clicking Edit opens the filter's property pages. Clicking Add either opens property pages for a new filter or starts a wizard, if you enabled the Use Add Wizard option. Since the wizard simply fills in the properties for you,

the next sections jump right in to the property pages. The Description page has only a single field for entering a free-form description of the filter. The following sections cover the other two pages, Addressing and Protocol.

Addressing Properties You use the Addressing page, shown in Figure 8-16, to specify the source and destination addresses you want the filter to match. The figure shows a specific IP Subnet selected for the Source address, but you have the following options:

- My IP Address is the address of the IPsec server.
- Any IP Address means that any IP address passes the filter.
- A Specific IP Address displays fields on the page for entering an IP address and subnet mask to use in the filter.
- A Specific IP Subnet also displays fields on the page, but you should only enter a subnet mask for use in the filter.

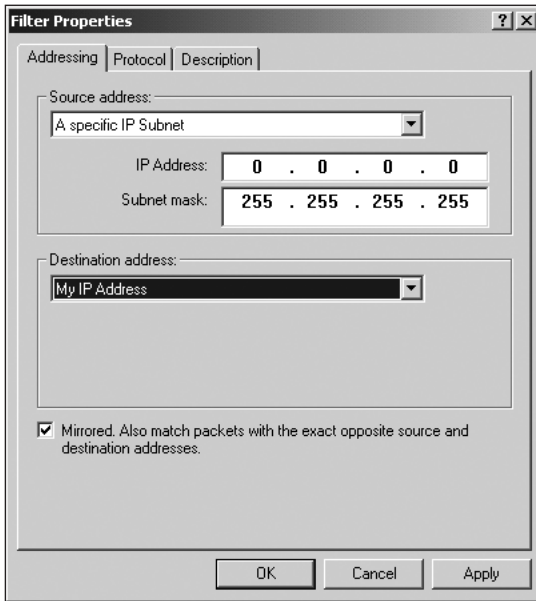


Figure 8-16 Addressing page for a filter

The Destination address presents all of the same options plus an option for specifying a particular DNS name. You use the source and destination addresses in combination in the filter. For example, you might want to create a filter that only matches hosts using the subnet 255.255.255.0 to connect to a specific IPsec server.

The Mirrored option, featured at the bottom of the Addressing page, makes a filter reciprocal. For example, choosing this option for the example in the previous paragraph creates a

filter that matches a specific IPSec server connecting to hosts on the subnet 255.255.255.0. This is useful for creating both inbound and outbound filters simultaneously.

Protocol Properties The Protocols page, shown in Figure 8-17, lets you match traffic being sent or received on a particular port or protocol. For example, you might want to match all ICMP traffic or all TCP traffic coming in over port 80.

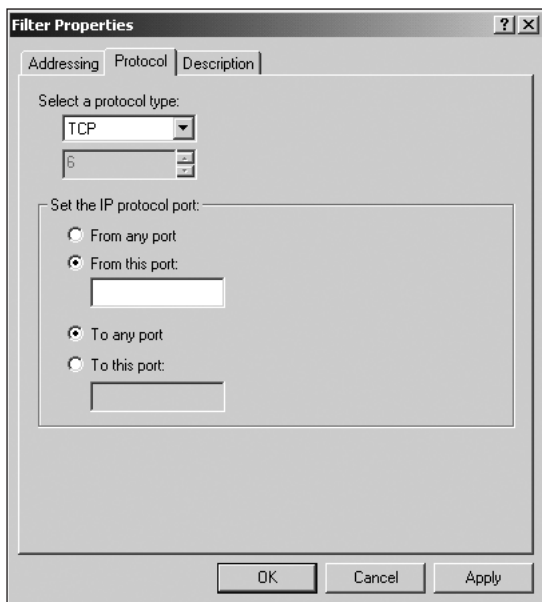


Figure 8-17 Protocols page for a filter

Managing Filter Actions

You use a filter list to match a connection and a filter action to define what happens when a match is made. As you saw when creating rules previously, a filter list and a filter action always work together to produce a desired result. The Manage Filter Actions page, shown in Figure 8-18, defines actions available to policies. By default, you get three actions: permit the connection, request security before allowing the connection, and require security before allowing the connection.

You can use the Edit button to edit properties for a selected action and the Add button either to open blank property dialog boxes or launch a wizard. The General page has only two parameters: they allow you to name and describe the action. The Security Methods page, shown in Figure 8-19, is where all the action happens.

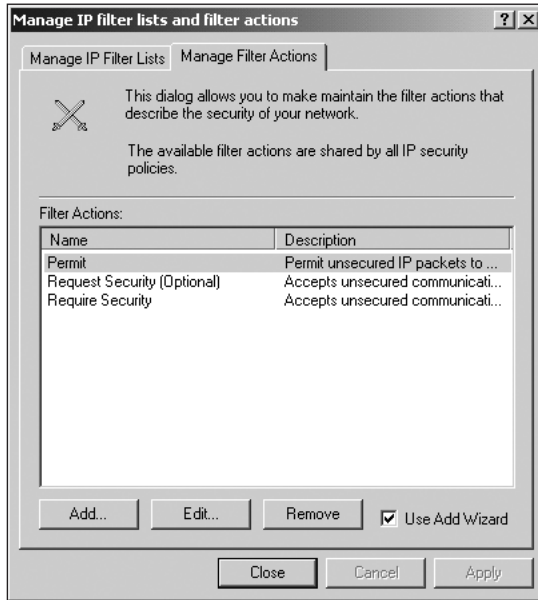


Figure 8-18 Managing Filter Actions

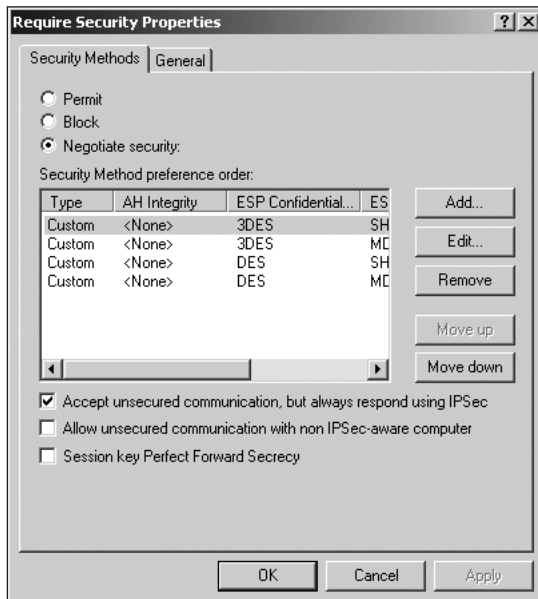


Figure 8-19 Security Methods of an action

This page presents three basic options that the action can perform when a connection matches a filter list: permit the connection with no further intervention, block the connection

altogether, or negotiate security for the connection. All of the remaining controls on the page relate to the Negotiate Security option and are unavailable when you select Permit or Block. These controls include:

- *Security Method preference order list*: shows which security methods the connection is allowed to use. You can add new methods or edit existing methods. Methods are attempted starting at the top of the list and working down, so you also have the option of moving methods up and down the list.
- *Accept unsecured communication, but always respond using IPSec* option: sets it up so that incoming requests are always answered by an attempt at an IPSec negotiation. If no IPSec connection can be made, the connection is allowed to proceed anyway.
- *Allow unsecured communication with non IPSec-aware computer*: lets computers not configured with IPSec make the connection anyway.
- *Session key Perfect Forward Secrecy*: prohibits the reuse of keying material.

Applying Policies to the Active Directory

Most of this chapter focuses on applying policies to a local computer, but knowing how to apply policies to the Active Directory is also valuable. First, you must configure the IPSec snap-in to configure default policies for a domain: either the local domain or a remote trusted domain. See Figure 8-1 for a refresher.

For the most part, policy management is exactly the same. You define policies, rules, authentications, filter lists, and filter actions. The difference comes when it is time to attach the policy to a domain or organizational unit within Active Directory. For this, you use the Group Policies snap-in, which you can add to a console in the same way you added the IPSec snap-in. Hands-on Project 8-3 at the end of the chapter outlines the steps for doing this.

IPSec policies must follow the same rules that apply to other objects assigned by group policy. Even though this chapter does not go into detail on Group Policy management, it presents four pretty simple rules to keep in mind:

- A policy applied at the domain level always overrides a policy applied at the local computer level.
- A policy applied to an organizational unit overrides policies applied at the domain level.
- If you have configured a hierarchy of organizational units, policies applied at lower levels in the hierarchy override policies applied at higher levels in the hierarchy.
- If you assign an IPSec policy and then delete the Group Policy object that created the policy, the policy remains in effect. The IPSec policy agent simply figures that the Group Policy object is unavailable and uses a cached version of the policy from the local computer. You must actually unassign the policy before removing the Group Policy object.

MANAGING AND MONITORING IPSEC

Once created and configured, policies have several management options related to them. You already learned how to assign and unassign a policy using the policy shortcut menu. In addition, you can use the shortcut menu to perform the following actions:

- *Check Policy Integrity command*: verifies that any changes you made to policy settings have been properly propagated by Group Policy to the computer accounts in the Group Policy Object (GPO). When you select the command, the IPsec snap-in just returns a dialog stating whether integrity is good or bad.
- *Restore Default Policies command*: restores all predefined default policies to their original state. This does not affect any new policies you create.
- *Import and Export Policies commands*: move policies between consoles so that you may copy policies to different local computers or domains once you create them.

The final tool for monitoring IPsec discussed in this chapter is actually named the IPsec Monitor. This very simple tool allows you to view the active security associations on local and remote computers. To activate the tool, use the Run command on the Start menu and issue the command `ipsecmon.exe`. If you want to manage a remote computer, you can add the computer's name after the command. Figure 8-20 shows a sample screen from IP Security Monitor.

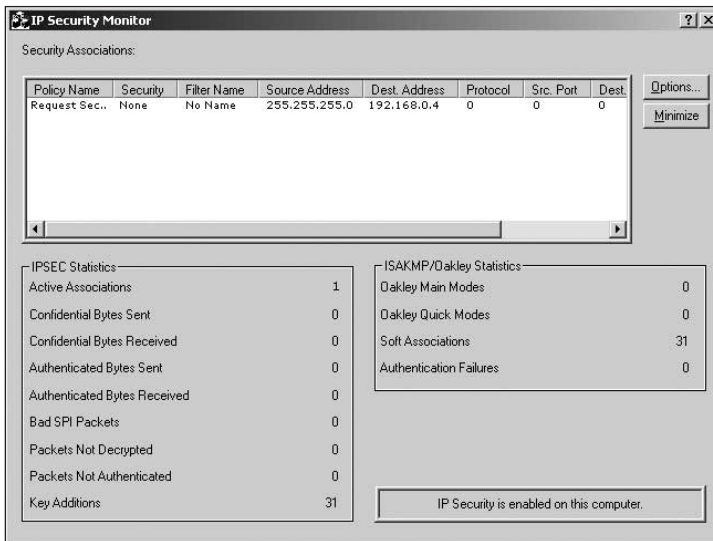


Figure 8-20 IPsec Monitor

There's really not much to configure for this tool. The only commands you can issue within it are to minimize the screen and open an options dialog box that lets you set the interval at which the monitor refreshes itself. The default refresh rate is every 15 seconds. As you can see, the monitor provides a good bit of information about each active security association, including the policy name, security level, filter name, and the source and destination addresses. In addition, the monitor shows you a number of statistics related to IPSec itself and to the IASKMP/Oakley Service.

CHAPTER SUMMARY

- IP Security (IPSec), an extension of the IP protocol, provides point-to-point authentication and encryption of data being sent between two computers on an IP-based network. Like IP, IPSec works at the Network layer. This means that higher-level protocols and applications in the TCP/IP protocol suite, like FTP, have nothing to do with the encryption process. They carry out their functions normally, passing data down the protocol layers, unaware of whether their data is eventually encrypted or not. Three forms of authentication (Kerberos, certificates, and pre-shared keys) are available to use with IPSec by default.
- IPSec can operate in two different modes, depending on the scope of the communication. These two modes of operation are transport mode and tunnel mode. In transport mode, two computers configured to use IPSec create a security association between themselves and carry out secure communication. In tunnel mode, an IPSec connection is created between two routers that connect two networks over a transit internetwork. Computers on one network can communicate with computers on the other network without having IPSec configured on the communicating computers.
- IPSec is actually installed by default on any Windows 2000 computer. All you do to enable it is to create an MMC console using the IP Security Management snap-in and then assign policies to be used. You can also use the snap-in to create and edit new policies. A policy is essentially a set of rules that governs a connection. Each rule is defined with a filter list, which is a list of filters a connection must pass to be considered a match; and a filter action, which is the action taken for any connection that makes the match.
- Once you create and assign policies, IPSec is ready to go. You can monitor it on a local or remote computer using the IPSec Monitor tool, which displays configured and active connections and a number of IPSec statistics.

KEY TERMS

- authentication** — A method for validating the identity of a user or a computer. IPSec supports three modes of encryption: Kerberos, certificates, and pre-shared keys.
- decrypt** — Process of decoding encrypted data.
- encrypt** — Process of sealing data using a special coding algorithm so that only intended recipients can decrypt and read it.
- filter action** — Actions assigned to a connection whose properties match an associated list of filters. Typical actions are to accept and block connections or to negotiate security for the connection.
- filter list** — List of filters assigned to a rule. Connections whose properties match the list of filters have an associated filter action applied to them.
- Internet Protocol (IP)** — Protocol in the TCP/IP protocol suite responsible for routing data over a network.
- IP Security (IPSec)** — Extension to the Internet Protocol (IP) used to secure data being sent between two computers on a network.
- IPSec client** — Computer that initiates the IPSec connection.
- IPSec driver** — IPSec component that actually encrypts and decrypts data using keys prepared by the ISAKMP/Oakley Service, and sends the data between computers.
- IPSec policies** — Sets of rules assigned to clients that define how those clients use IPSec.
- IPSec policy agent service** — IPSec component responsible for retrieving the computer's assigned IPSec policy from the Active Directory.
- IPSec server** — Computer that responds to an IPSec connection.
- ISAKMP/Oakley Service** — IPSec component that creates the security association between communicating computers and is also responsible for generating the keys used to encrypt and decrypt the data sent over the IPSec connection.
- Kerberos V5** — Default authentication system used by Windows 2000. It is an open standard widely-supported by other operating systems, as well.
- pre-shared keys** — Passwords entered into each computer communicating with IPSec. As long as both computers are configured with the same pre-shared key, they trust one another.
- public key certificates** — Provided by a certificate authority. Each end of the IPSec connection uses the other end's public certificate for authentication.
- security association** — Defines the common security mechanisms, such as keys, that two computers use to create the IPSec connection.
- transport mode** — Mode in which the two endpoints of IPSec communication are two computers that have IPSec configured. For this mode to work, both computers must use the TCP/IP protocol.
- tunnel mode** — Mode in which two communicating computers do not use IPSec themselves. Instead, the gateways connecting each client's LAN to the transit network create a virtual tunnel that uses the IPSec protocol to secure all communication that passes through it.

REVIEW QUESTIONS

1. At what level of the OSI networking model does IPSec work?
 - a. Application layer
 - b. Transport layer
 - c. Network layer
 - d. Physical layer
2. Which of the following is a responsibility of the IPSec policy agent?
 - a. Retrieves policy information from the Active Directory
 - b. Generates keys based on defined policies
 - c. Oversees the creation of a security association
 - d. Encrypts and decrypts data based on security keys
3. Which of the following is *not* an available form of IPSec authentication.
 - a. Kerberos
 - b. Pre-shared keys
 - c. Windows Integrated
 - d. Clear Text
 - e. Certificates
4. The _____ is the IPSec component responsible for creating the keys used to encrypt and decrypt data.
5. The ISAKMP/Oakley Service is responsible for creating a security association. True or false?
6. Which of the following networking protocols may be used by two computers that communicate between networks configured to use IPSec in tunnel mode?
 - a. TCP/IP
 - b. IPX/SPX
 - c. AppleTalk
 - d. All of the above
7. Restarting the IPSec driver is the best way to restart the IPSec Policy Agent. True or false?
8. What must you do to enable IPSec on a local computer?
 - a. Install IPSec using the Add/Remove Software Control Panel applet.
 - b. Create an IPSec policy using the IPSec snap-in.
 - c. Assign an IPSec policy using the IPSec snap-in.
 - d. Install the IPSec Policy Agent.

9. Which of the following can you use to verify that IPSec is running on a local computer?
 - a. The ipconfig tool
 - b. The ipsecmon tool
 - c. The tracert tool
 - d. All of the above
10. When defining a rule, you must associate a filter list with a filter action. True or false?
11. The _____ protocol is used in conjunction with IPSec to create a Virtual Private Network.
12. You are configuring IPSec in tunnel mode between two remote networks. How many total rules do you need to configure?
 - a. 1
 - b. 2
 - c. 4
 - d. 8
13. You use the Perfect Forward Secrecy options available for sessions and rules to make sure that _____.
14. Which of the following filter actions are available by default? (Choose three.)
 - a. Permit a connection.
 - b. Block a connection.
 - c. Request security before allowing the connection.
 - d. Require security before allowing the connection.
15. A _____ is a set of rules governing how a client uses IPSec when making a connection.
16. In what order are IPSec policies applied?
 - a. Local, domain, then organizational unit
 - b. Organizational unit, domain, then local
 - c. Domain, local, then organizational unit
 - d. Domain, organizational unit, then local
17. In one form of IPSec authentication, _____, identical passwords are entered into each computer communicating with IPSec.
18. Restoring default IPSec policies in the IPSec snap-in deletes any custom policies and restores the original policies to their default configuration. True or false?

19. Which of the following is true of IPSec authentication?
 - a. It only lets you enable one mode of authentication per rule.
 - b. It allows multiple authentications, but always uses Kerberos as the default method.
 - c. It allows multiple authentications, but always uses pre-shared keys as the default method.
 - d. It prevents you from using Kerberos and certificates together but lets you combine either of these with pre-shared keys.
20. When configuring IPSec in tunnel mode, only DNS names can specify tunnel endpoints. True or false?

HANDS-ON PROJECTS

All Hands-on Projects in this chapter require at least one server computer set up as described in the lab set-up section in the front of this book.



Project 8-1

To enable IPSec on a local computer:

1. Click **Start** and then click **Run**.
2. In the **Run** field, type **mmc** and click **OK**.
3. From the **Console** menu of the Microsoft Management Console, select the **Add/Remove snap-in** command.
4. Click the **Add** button.
5. From the **Available Standalone Snap-Ins** list, select the **IP Security Policy Management** entry and click **Add**.
6. Make sure that the **Local Computer** option is selected, and click **Finish**.
7. Click **Close** to close the **Add Standalone Snap-In** dialog box.
8. Click **OK** to close the **Add/Remove Snap-In** dialog box.
9. In the left pane of the MMC main window, select the **IP Security Policies on Local Machine** object.
10. In the policies list in the right pane, right-click the **Server (Request Security)** policy and select the **Assign** command.
11. Verify that the entry in the Policy Assigned column for that policy changed to **Yes**.



Project 8-2

This project assumes that the MMC console that you created in Hands-on Project 8-1 is still open. If not, you must open or recreate it to proceed.

To configure the properties for a policy:

1. In the left pane of the MMC main window, select the **IP Security Policies on Local Machine** object.
2. In the policies list in the right pane, right-click the **Server (Request Security)** policy and select the **Properties** command.
3. Select the **All IP Traffic** rule, and then click the **Edit** button.
Note that this rule matches all IP packets from the local computer to any other computer.
4. Click the **Filter Action** tab.
5. Change the default setting (**Request Security**) to **Require Security**.
6. Click the **Edit** button.
7. Disable the **Accept unsecured communication, but always respond using IPSec** option.
8. Click **OK** to return to the **Edit Rule Properties** dialog box.
9. Click the **Authentication Methods** tab.
10. Click the **Add** button.
11. Select the **Use this string to protect the key exchange** option.
12. In the field below the option, type a **password**.
13. Click **OK** to return to the **Edit Rule Properties** dialog box.
14. Click **Close** to return to the Policy property pages.
15. Click **Close** to return to the IPSec snap-in.



Project 8-3

To enable IPSec for an entire domain:

1. Click **Start** and then click **Run**.
2. In the **Run** field, type **mmc.exe**, and click **OK**.
3. From the **Console** menu of the Microsoft Management Console, select the **Add/Remove snap-in** command.
4. Click the **Add** button.
5. From the **Available Standalone Snap-Ins** list, select the **Group Policy** entry and click **Add**. The **Select Group Policy Object** dialog box opens.

6. Click the **Browse** button to open a dialog box that lets you specify a group policy object.
7. Select the **Default Domain Policy** entry, and click **OK**.
8. Click **Finish** to return to the **Add Standalone Snap-In** dialog box.
9. Click the **Close** button to return to the **Add/Remove Snap-In** dialog box.
10. Click the **OK** button to return to the console.
11. Expand the **Default Domain Policy** object in the left pane of the console until you can find and then select the **IP Security Policies on Active Directory** object. The right pane shows available policies in the Active Directory.
12. Right-click the **Server (Request Security)** policy, and select **Assign** from the shortcut menu.



Project 8-4

To create the local end of an IPSec tunnel:

1. Right-click the **IP Security Policies on Local Machine** object, and select the **Create IP Security Policy** command to open the New Policy Wizard.
2. Click **Next** to skip the Welcome window.
3. Enter a **name** for the new policy, and click the **Next** button to continue.
4. On the **Requests For Secure Communications** page, disable the **Activate Default Response Rule** option and click **Next**.
5. On the **Summary** page, make sure the **Edit Properties** button is selected and click **Finish**.
The property pages for the new policy open.
6. On the **Rules** tab, make sure the **Use Add Wizard** button is enabled and click the **Add** button.
7. Click **Next** to skip the wizard's Welcome window.
8. On the **Tunnel Endpoint** page, select the **Tunnel Endpoint is specified by this IP address** option, enter the **IP address** for the remote tunnel interface on the transit network, and click **Next** to continue.
9. On the **Network Type** page, select the **Local Area Network (LAN)** option and then click **Next**.
10. On the **Authentication Method** page, select the **Windows 2000 Default (Kerberos V5)** option and then click **Next**.
11. On the **IP Filter List** page, select the **All IP Traffic** list and click **Next**.
12. On the **Filter Action** page, select the **Request Security (Optional)** action and click **Next**.
13. Click **Finish** to return to the policy properties dialog box.
14. Click **Close** to return to the IPSec snap-in.

CASE PROJECTS



Case 1

You are the network administrator for a large company based in San Francisco. Your company has just acquired a smaller company in Boston and you have been given the task of joining the two networks. You plan to connect the networks to one another using the Internet. The network in San Francisco is TCP/IP-based and the one in Boston is IPX/SPX-based. Each is connected to the Internet via a router to an Internet Service Provider. You would like to establish a Virtual Private Network between the two networks and secure it using IPSec. Write out a plan for this. Include what mode you would use IPSec in and whether any additional protocols would be needed on the networks.



Case 2

You have now finished planning the IPSec configuration from Case 1 and have enabled IPSec on one server from each network. These servers will be governing the IPSec communications over the routers. Describe the rules you would need to put in place on each end of the connection so that the IPSec configuration is complete.

